

Nested Satisfiability

by Donald E. Knuth*

Abstract. A special case of the satisfiability problem, in which the clauses have a hierarchical structure, is shown to be solvable in linear time, assuming that the clauses have been represented in a convenient way.

Let X be a finite alphabet linearly ordered by $<$; we will think of the elements of X as boolean variables. As usual, we define the *literals* over X to be elements of the form x or \bar{x} , where $x \in X$. Literals that belong to X are called positive; the others are called negative.

The linear ordering of X can be extended to a linear preordering of all its literals in a natural way if we simply disregard the signs. For example, if $X = \{a, b, c\}$ has the usual alphabetic order, we have

$$a \equiv \bar{a} < b \equiv \bar{b} < c \equiv \bar{c}.$$

If σ and τ are literals, we write $\sigma \leq \tau$ if $\sigma < \tau$ or $\sigma \equiv \tau$; this holds if and only if the relation $\sigma > \tau$ is false.

A *clause* over X is a set of literals on distinct variables. Thus, the literals of a clause can be written in increasing order,

$$\sigma_1 < \sigma_2 < \cdots < \sigma_k.$$

A set \mathcal{C} of clauses over X is *satisfiable* if there exists a clause over X that has a nonempty intersection with every clause in \mathcal{C} . For example, the clauses

$$\{a, \bar{b}, c\} \quad \{\bar{a}, \bar{c}\} \quad \{\bar{a}, b, c\} \quad \{\bar{a}, \bar{c}\} \quad \{a, b\}$$

over $\{a, b, c\}$ are satisfiable uniquely by the clause $\{a, b, \bar{c}\}$.

We say that clause \mathcal{C} *straddles* clause \mathcal{C}' if there are literals σ, τ in \mathcal{C} and ξ' in \mathcal{C}' such that

$$\sigma < \xi' < \tau.$$

Two clauses *overlap* if they straddle each other. For example, $\{a, \bar{b}, c\}$ and $\{\bar{a}, b, c\}$ overlap; but the other nine pairs of clauses in the example above are non-overlapping. Clauses on two elements each, like $\{a, c\}$ and $\{b, d\}$, can also be overlapping. A set of clauses in which no two overlap is called *nested*.

The general problem of deciding whether a given set of clauses is satisfiable is well known to be NP-complete. But we will see that the analogous question for nested clauses is efficiently decidable. The main reason for interest in nested clauses is David Lichtenstein's theorem of *planar satisfiability*, which can be restated in algebraic terms as follows: The

* Computer Science Department, Stanford University; research supported in part by National Science Foundation grant CCR-8610181.

joint satisfiability problem for *two* sets $\mathcal{C}, \mathcal{C}'$ of nested clauses is NP-complete. In fact, Lichtenstein proved [1, p. 339] that this problem is NP-complete even if all clauses of \mathcal{C} contain only positive literals and all clauses of \mathcal{C}' contain only negative literals, with at most three literals per clause.

1. Structure of nested clauses. A clause over an ordered alphabet has a least literal σ and a greatest literal τ . Any variable that lies strictly between σ and τ is said to be *interior* to that clause. A variable can occur as an interior literal at most once in a set of nested clauses; for if it is an interior literal in two different clauses, those clauses overlap. Hence, the total number of elements among m nested clauses on n variables is at most $2m + n$.

Let us write $\mathcal{C} \succ \mathcal{C}'$ if \mathcal{C} straddles \mathcal{C}' but \mathcal{C}' does not straddle \mathcal{C} . This relation is transitive. For if $\mathcal{C} \succ \mathcal{C}'$ and \mathcal{C}' straddles \mathcal{C}'' , we have literals

$$\sigma < \xi' < \tau, \quad \sigma' < \xi'' < \tau'$$

in appropriate clauses; and we must have $\sigma \leq \sigma'$ and $\tau' \leq \tau$, or else \mathcal{C}' would straddle \mathcal{C} . Hence \mathcal{C} straddles \mathcal{C}'' . Similarly if $\mathcal{C}' \succ \mathcal{C}''$ and \mathcal{C}'' straddles \mathcal{C} , then \mathcal{C}' straddles \mathcal{C} . Therefore $\mathcal{C} \succ \mathcal{C}' \succ \mathcal{C}''$ implies that $\mathcal{C} \succ \mathcal{C}''$.

In a set of nested clauses, we have $\mathcal{C} \succ \mathcal{C}'$ if and only if \mathcal{C} straddles \mathcal{C}' . The transitivity of this relation implies that we can topologically sort any set of nested clauses into a linear arrangement in which each clause appears after every clause it straddles. When such an arrangement is given, and when the elements of each clause are presented in order, we will show that satisfiability can be decided in $O(m + n)$ steps on a RAM, where m is the number of clauses and n is the number of variables.

(Incidentally, a set of nested clauses can be shown to have a tree-like structure, although we do not need this characterization in the algorithm. Let us write $\mathcal{C} \leq \mathcal{C}'$ if $\sigma \leq \tau'$ for all $\sigma \in \mathcal{C}$ and $\tau' \in \mathcal{C}'$. If neither \mathcal{C} nor \mathcal{C}' straddles the other, it is easy to see that we must have either $\mathcal{C} \leq \mathcal{C}'$ or $\mathcal{C}' \leq \mathcal{C}$, unless \mathcal{C} and \mathcal{C}' are both clauses on the same two literals. Suppose we call such 2-element clauses equivalent. Then a nested set of clauses will satisfy the condition

$$(\mathcal{C} \succ \mathcal{C}'' \text{ and } \mathcal{C}' \succ \mathcal{C}'') \text{ implies } (\mathcal{C} \succ \mathcal{C}' \text{ or } \mathcal{C} \equiv \mathcal{C}' \text{ or } \mathcal{C}' \succ \mathcal{C}),$$

because we cannot have $\mathcal{C} \succ \mathcal{C}''$ and $\mathcal{C}' \succ \mathcal{C}''$ when $\mathcal{C} \leq \mathcal{C}'$. This means that \succ is the ancestor relation in a hierarchy.)

2. An algorithm. Let us assume that the alphabet X is represented as the positive integers $\{1, 2, \dots, n\}$, with $\bar{x} = -x$. The clauses will be specified in two arrays

$$lit[1 \dots 2m + n] \quad \text{and} \quad start[1 \dots m + 1]$$

where the literals of clause i are

$$lit[j], \quad \text{for } start[i] \leq j < start[i + 1]$$

in increasing order as j increases. The clauses are assumed to be arranged so that clause i does not straddle clause i' when $i < i'$. We can safely assume that all clauses contain at least two literals.

The key idea of the algorithm below is that the interior variables of a clause are not present in subsequent clauses. Therefore we only need to remember information about the dynamically changing set of all variables

$$1 = x_1 < x_2 < \cdots < x_k = n$$

that have not yet appeared as interior variables. Initially $k = n$.

The set of all clauses seen so far, as the algorithm proceeds to consider the clauses in turn, can be conceptually partitioned into intervals

$$[x_1 \dots x_2], [x_2 \dots x_3], \dots, [x_{k-1} \dots x_k],$$

such that all literals of each previously processed clause belong to one of these intervals. The current intervals are maintained in an array

$$next[1 \dots n]$$

where $next[x_j] = x_{j+1}$ for $1 \leq j < k$.

The only slightly complex data structure in the algorithm below is the array

$$sat[1 \dots n, boolean, boolean]$$

which has the following interpretation: If $[x_j \dots x_{j+1}]$ is an interval of the current partition, then $sat[x_j, s, t]$ will be either 0 or 1 for each pair $s, t \in \{false, true\}$. It is 1 if and only if the clauses already processed, belonging to the interval $[x_j \dots x_{j+1}]$, are satisfiable by clauses in which the least and greatest literals are respectively $x_j|s$ and $x_{j+1}|t$, where

$$x|s = \begin{cases} -x, & s = false; \\ +x, & s = true. \end{cases}$$

For example, suppose we have seen only one clause, $\{1, -2\}$. Then we will have

$$\begin{aligned} sat[1, false, true] &= 0; \\ sat[1, false, false] &= sat[1, true, false] = sat[1, true, true] = 1. \end{aligned}$$

It turns out that the sat array contains all the information necessary to continue processing, because literals that have appeared as interior variables will not be present in subsequent clauses.

The algorithm's main task is to maintain the sat array as it examines a new clause $\mathcal{C}_i = \{\sigma_1, \dots, \sigma_q\}$. The variables $|\sigma_1| < \cdots < |\sigma_q|$ will be a subset of the current partition variables x_1, \dots, x_k . All of the current partition variables between $|\sigma_1|$ and $|\sigma_q|$, whether they appear in the new clause or not, are interior to the clause, so they will be removed.

Suppose $|\sigma_1| = x_p$. The algorithm proceeds by letting a variable x run through the values $x_p, x_{p+1}, \dots, |\sigma_q|$, maintaining information needed to update the values of $sat[x_p, s, t]$ when the interior variables of C_i are eliminated from the partition. Let $C_i(x)$ be the literals of C_i that are strictly less than x , and let $\mathcal{C}(x)$ be the clauses preceding C_i whose literals are confined to the interval $[x_p \dots x]$. The updating process is carried out by computing auxiliary values $newsat_x[s, t]$ defined as follows:

$$newsat_x[s, t] = \begin{cases} 0, & \text{if } \mathcal{C}(x) \text{ is not satisfiable}(s, t); \\ 1, & \text{if } \mathcal{C}(x) \text{ is satisfiable}(s, t) \text{ but } \mathcal{C}(x) \cup \{C_i(x)\} \text{ isn't}; \\ 2, & \text{if } \mathcal{C}(x) \cup \{C_i(x)\} \text{ is satisfiable}(s, t). \end{cases}$$

Here ‘satisfiable(s, t)’ means there is a clause containing $x_p|s$ and $x|t$ that has a nonempty intersection with each clause of the given set of clauses.

For example, suppose $C_i = \{-1, 2, 4\}$ and $\{x_1, x_2, x_3, x_4\} = \{1, 2, 3, 4\}$, and suppose that the clauses $\mathcal{C}_1, \dots, \mathcal{C}_{i-1}$ have led to the following values:

s	t	$sat[1, s, t]$	$sat[2, s, t]$	$sat[3, s, t]$
<i>false</i>	<i>false</i>	0	0	0
<i>false</i>	<i>true</i>	1	1	0
<i>true</i>	<i>false</i>	1	1	0
<i>true</i>	<i>true</i>	1	0	1

Then we have

s	t	$newsat_1[s, t]$	$newsat_2[s, t]$	$newsat_3[s, t]$	$newsat_4[s, t]$
<i>false</i>	<i>false</i>	1	0	2	0
<i>false</i>	<i>true</i>	0	2	0	0
<i>true</i>	<i>false</i>	0	1	2	0
<i>true</i>	<i>true</i>	1	1	1	1

and we will want to update the arrays by setting $next[1] \leftarrow 4$ and

```

 $sat[1, false, false] \leftarrow 0;$ 
 $sat[1, false, true] \leftarrow 0;$ 
 $sat[1, true, false] \leftarrow 0;$ 
 $sat[1, true, true] \leftarrow 1.$ 

```

If C_i were $\{-1, 2, -4\}$ instead of $\{-1, 2, 4\}$, the computation of $newsat$ would be the same, but the values of $sat[1, s, t]$ would all become 0; the clauses would be unsatisfiable, since $newsat_4[true, true]$ is only 1, not 2. (The reader is encouraged to study this example carefully, because it reveals the key principles underlying the algorithm.)

3. Programming details. It is convenient to assume that an artificial $(m + 1)$ st clause with the dummy variables $\{0, n + 1\}$ has been added after C_m . Therefore we will declare slightly larger arrays than stated earlier:

```
 $start[1 \dots m + 2]; next[0 \dots n]; sat[0 \dots n, boolean, boolean].$ 
```

There are two auxiliary arrays $newsat[boolean, boolean]$ and $tmp[boolean, boolean]$. We can now decide the nested satisfiability problem as follows.

```

for  $x \leftarrow 0$  to  $n$  do  $next[x] \leftarrow x + 1$ ;
for  $x \leftarrow 0$  to  $n$  do
  for  $s \leftarrow \text{false}$  to  $\text{true}$  do for  $t \leftarrow \text{false}$  to  $\text{true}$  do  $sat[x, s, t] \leftarrow 1$ ;
for  $i \leftarrow 1$  to  $m + 1$  do
  begin  $l \leftarrow \text{abs}(\text{lit}[start[i]])$ ;  $r = \text{abs}(\text{lit}[start[i + 1] - 1])$ ;
  ⟨Compute the  $newsat$  table⟩;
   $next[l] \leftarrow r$ ;
  for  $s \leftarrow \text{false}$  to  $\text{true}$  do for  $t = \text{false}$  to  $\text{true}$  do
     $sat[l, s, t] \leftarrow newsat[s, t] \text{ div } 2$ ;
  end;
if  $sat[0, \text{true}, \text{true}] = 1$  then print (‘Satisfiable’) else print (‘Unsatisfiable’).

```

The example in the previous section illustrates how the $newsat$ table can be computed in general. We run the process slightly longer so that a good $newsat$ value will be 2 (not 1) at the end. (The value of σ_q must be examined.)

⟨Compute the $newsat$ table⟩ =

```

 $j \leftarrow start[i]$ ;  $sig \leftarrow \text{lit}[j]$ ;  $x \leftarrow \text{abs}(sig)$ ;
 $newsat[\text{false}, \text{false}] \leftarrow 1$ ;  $newsat[\text{true}, \text{true}] \leftarrow 1$ ;
 $newsat[\text{false}, \text{true}] \leftarrow 0$ ;  $newsat[\text{true}, \text{false}] \leftarrow 0$ ;
while  $\text{true}$  do
  begin if  $x = \text{abs}(sig)$  then
    begin ⟨Upgrade a  $newsat$  from 1 to 2, if possible⟩;
     $j \leftarrow j + 1$ ;  $sig \leftarrow \text{lit}[j]$ ;
    if  $j = start[i + 1]$  then goto done;
    end;
  ⟨Modify  $newsat$  for the next  $x$  value⟩;
   $x \leftarrow next[x]$ ;
  end;
done:

```

⟨Upgrade a $newsat$ from 1 to 2, if possible⟩ =

```

 $t \leftarrow (x = sig)$ ;
for  $s \leftarrow \text{false}$  to  $\text{true}$  do
  if  $newsat[s, t] = 1$  then  $newsat[s, t] \leftarrow 2$ .

```

⟨Modify $newsat$ for the next x value⟩ =

```

for  $s \leftarrow \text{false}$  to  $\text{true}$  do for  $t \leftarrow \text{false}$  to  $\text{true}$  do
   $tmp[s, t] \leftarrow \max(newsat[s, \text{false}] * sat[x, \text{false}, t],$ 
     $newsat[s, \text{true}] * sat[x, \text{true}, t])$ ;
for  $s \leftarrow \text{false}$  to  $\text{true}$  do for  $t \leftarrow \text{false}$  to  $\text{true}$  do
   $newsat[s, t] \leftarrow tmp[s, t]$ .

```

The running time is $O(m + n)$, because each value of x is either first or last in the current clause (accounting for $2(m + 1)$ cases) or it is being permanently removed from the partition (in exactly n cases, because of the dummy clause $\{0, n + 1\}$ at the end).

We have not considered here the time that might be required to test if a given satisfiability problem is, in fact, nested under some ordering of its variables.

Concluding remarks. This algorithm for nested satisfiability works by essentially replacing each clause by a clause containing only two literals, using a special form of “dynamic 2SAT” to justify the replacement. However, the instances of 2SAT that arise are not completely general. This suggests that a somewhat larger special case of the satisfiability problem might be solvable in linear time by similar techniques.

Acknowledgment. I wish to thank Andrew Goldberg for posing the problem of nested satisfiability during a conversation about Lichtenstein’s theorem, and I wish to thank the referees for their helpful remarks.

References

- [1] Lichtenstein, D.: Planar Formulæ and Their Uses. *SIAM J. Comput.* **11**, 329–343 (1982)